

Resources Request Examples

Category: PBS on Pleiades

Since Pleiades consists of three different processor types, Harpertown, Nehalem-EP and Westmere-EP, keep the following in mind when requesting PBS resources for your job:

- Starting May 1, 2011, charging on the usage of the three Pleiades processor types is based on a common Standard Billing Unit which is on a per-node basis. The SBU rate for each of the Pleiades processor type is:

Processor Type	SBU Rate (per node)
Westmere-EP	1 (12 cores in a node)
Nehalem-EP	0.8 (8 cores in a node)
Harpertown	0.45 (8 cores in a node)

The actual amount of memory per node through PBS is slightly less, 7.6 GB/node for Harpertown and 22.5 GB/node for Nehalem-EP and Westmere-EP.

Use the "model=[har,neh,wes]" attribute to request the processor type(s) for your job. If the processor type is not specified in user's PBS resource list, the job is routed to use the default processor type, Harpertown.

Example 1:

Here are some examples of requesting certain processor models for a 128-process job:

```
#PBS -l select=16:ncpus=8:model=har
# to run all 8 cores on each of 16 Harpertown nodes

#PBS -l select=32:ncpus=4:model=har
# to run on only 4 cores on each of 32 Harpertown nodes
# (note: will be charged for 32 nodes = 256 cores)

#PBS -l select=16:ncpus=8:model=neh
# to run all 8 cores on each of 16 Nehalem-EP nodes

#PBS -l select=11:ncpus=12:model=wes
# to run all 12 cores on each of 11 Westmere-EP nodes
# (4 cores in 11th node will go unused)
```

Note that you can specify both the queue type (-q normal, debug, long, wide, low) and the processor type (-l model=har,neh,wes). For example:

```
#PBS -q normal
```

```
#PBS -l select=16:ncpus=8:model=neh
```

If your application can run on any of the three processor types, you may want to submit your job to a processor type that has more unoccupied nodes by other running jobs. This can possibly reduce the wait time of your job. The script *node_stats.sh* provides information about the total, used and free nodes for each processor type. For example:

```
%/u/scicon/tools/bin/node_stats.sh
```

```
Pleiades Nodes Total: 9394
Pleiades Nodes Used : 8128
Pleiades Nodes Free : 1266
```

```
Harpertown Total: 5854 Used: 4878 Free: 976
Nehalem      Total: 1255 Used: 1036 Free: 219
Westmere     Total: 2285 Used: 2214 Free: 71
```

```
Currently queued jobs are requesting: 1463 Harpertown, 1767 Nehalem,
2860 Westmere nodes
```

Add `"/u/scicon/tools/bin"` to your PATH in `.cshrc` or other shell start up scripts to avoid having to type in the complete path for this tool.

You can also find for each job what processor models are used by looking at the "Model" field of the output of the command:

```
%qstat -a -W o=+model
```

- The Harpertown nodes in rack 32 have 16 GB memory/node instead of 8 GB/node.

Example 2:

This example shows a request of 2 nodes with bigmem in rack 32.

```
#PBS -l select=2:ncpus=8:bigmem=true:model=har
```

- For a multi-node PBS job, the NCPUs used in each node can be different. This is useful for jobs that need more memory for some processes, but less for other processes. Resource requests can be done in "chunks" for a job with varying NCPUs per node.

Example 3:

This example shows a request of two chunks of resources. In the first chunk, 1 CPU in 1 node, and in the second chunk, 8 CPUs in each of three other nodes are requested:

```
#PBS -l select=1:ncpus=1+3:ncpus=8
```

- A PBS job can run across different processor types. This can be useful in two scenarios:
 1. when you can not find enough free nodes within one model for your job
 2. when some of your processes need more memory while others need much less

This can be accomplished by specifying the resources in "chunks", with one chunk asking for one processor type while another chunk asking for a different processor type.

Example 4

Here is an example to request 1 Westmere node (which provides 24 GB/node) and 3 Harpertown nodes (which provides 8 GB/node).

```
#PBS -lplace=scatter:excl:group=model  
#PBS -lselect=1:ncpus=12:mpiprocs=12:model=wes+3:ncpus=8:mpiprocs=8:model=har
```

Article ID: 188

Last updated: 23 May, 2011

Computing at NAS -> Running Jobs with PBS -> PBS on Pleiades -> Resources Request Examples

<http://www.nas.nasa.gov/hecc/support/kb/entry/188/?ajax=1>